

ABSTRACT

Text features corresponding to pieces of media content (e.g., images, audio, multimedia content, etc.) are extracted from media content sources. One or more text features (e.g., one or more words) for a piece of media content are extracted from text associated with the piece of media content and text feature vectors generated therefrom and used during subsequent searching. Additional low-level feature vectors may also be extracted from the piece of media content and used during the subsequent searching. Relevance feedback can also be received from a user(s) identifying the relevance of pieces of media content rendered to the user in response to his or her search request. The relevance feedback is logged and can be used in determining how to respond to subsequent search requests, such as by modifying feature vectors (e.g., text feature vectors) corresponding to the pieces of media content for which relevance feedback is received.